proceedings

MDPI

Proceeding Paper

Best Practices for the Responsible Adoption of Generative AI in Higher Education †

Alexander Mikroyannidis *, Audrey Ekuban, Joseph Kwarteng and John Domingue

Knowledge Media Institute, The Open University, Milton Keynes MK7 6AA, UK
* Correspondence: alexander.mikroyannidis@open.ac.uk
† Presented at the Online Workshop on Adaptive Education: Harnessing AI for Academic Progress, Online, 12 April 2024.

Abstract: In this paper, we propose a set of best practices for the responsible adoption of Generative Artificial Intelligence (AI) in higher education. These best practices provide a comprehensive framework for higher education institutions to effectively and ethically integrate Generative AI into their teaching and learning practices. The framework prioritises a responsible and human-centred approach, alongside pedagogical soundness, careful planning, transparency, as well as content quality. By exploring the responsible adoption of Generative AI in higher education, we seek to provide scalable, personalised learning experiences for large cohorts of students. Our research focuses on harnessing Generative AI to offer tailored educational content and generate constructive feedback for students. Additionally, by adhering to responsible AI practices, we aim to address challenges such as misinformation, copyright violations, and bias.

Keywords: generative AI; responsible AI; personalised learning; higher education

Academic Editors: Juliana Gerard and Muskaan Singh

Published: 6 March 2025

Citation: Mikroyannidis, A.; Ekuban, A.; Kwarteng, J.; Domingue, J. Best Practices for the Responsible Adoption of Generative AI in Higher Education. *Proceedings* **2025**, *114*, 6. https://doi.org/10.3390/proceedings2025114006

1. Introduction

For decades, educators have drawn inspiration from Bloom's study [1], which found that students receiving one-on-one instruction outperformed those in traditional classroom settings by two standard deviations. On average, individually taught students surpassed 98% of their peers in the control group. However, access to personalised education remains constrained by socioeconomic factors, making it a significant issue for educational equity. The rise of Generative AI presents new opportunities to deliver cost-effective personalised learning, significantly benefiting both students and teachers. By enhancing learning outcomes and expanding accessibility, this technology has the potential to support a more diverse range of learners.

UNESCO has published a comprehensive report on the capabilities and transformative potential of Generative AI in higher education [2]. Sharples [3] categorises the various educational roles that Generative AI can play in promoting collaborative and social learning. Additionally, researchers have examined the challenges associated with integrating Generative AI into higher education [4–6]. Key concerns include ethical considerations such as the potential for misuse, the creation of misleading or harmful content, and issues related to misinformation, copyright infringement, and algorithmic bias [7,8]. As the adoption of Generative AI tools accelerates among students and educators, higher education institutions must address these challenges to ensure the responsible and ethical use of AI-generated content [7].

On the other hand, the various challenges that the adoption of Generative AI poses for higher education have also been explored by researchers [4–6]. These studies point out

that Generative AI comes with ethical considerations, such as the potential for misuse, or the generation of deceptive or malicious content [7,8]. Issues related to misinformation, copyright violations, and algorithmic biases, as well as the responsible use of AI-generated content, need to be addressed by higher education institutions, especially since the use of Generative AI tools by students and educators is rapidly increasing [7].

This paper aims to address the above challenges by proposing a framework of best practices for the responsible adoption of Generative AI in higher education. This framework is based on our ongoing work, which leverages Generative AI to enhance both our teaching practices and the learning experiences of our students [9–12]. The remainder of this paper is structured as follows. First, we present the research project Smart Assessment and Guided Education with Responsible AI (SAGE-RAI), which aims to leverage Generative AI to enable cost-effective and scalable personalised education. We then present our approach for the ingestion of educational content by Generative AI. Finally, we propose a set of best practices for the responsible adoption of Generative AI in higher education, and we conclude this paper.

2. Smart Assessment and Guided Education with Responsible AI

The SAGE-RAI (Smart Assessment and Guided Education with Responsible AI) project, funded by the UKRI, explores the potential of Generative AI to provide cost-effective and scalable personalised education. This initiative aims to develop a platform for student assessment and guidance while responsibly leveraging Generative AI to tackle key lifelong learning challenges. These include the digital divide, which excludes learners without access to technology or digital literacy; financial barriers, which make accredited courses unaffordable for many; and bias in recognising prior learning, which often favours traditional education over informal or marginalised learning experiences. Additionally, a lack of inclusive pedagogy and culturally relevant content can limit the accessibility and relatability of lifelong learning courses. By addressing these challenges, the project seeks to make education more inclusive, equitable, and effective for diverse learners.

To achieve this mission, we are developing AI-powered digital assistant plugins that integrate seamlessly into existing online learning platforms. These are deployed in the Open Data Institute's (ODI) learning environment and OpenLearn, the Open University's free learning resource. We are also evaluating the educational benefits of Generative AI to ensure it enhances learning outcomes while maintaining student satisfaction. Additionally, we will establish a business case and process models to help educational providers adopt AI-driven tools efficiently. Ethical considerations are central to the SAGE-RAI project, with a focus on responsible AI implementation. To ensure fairness, transparency, and accuracy, fact-checking mechanisms, regular audits, and ethical reviews will monitor potential biases and misinformation in AI-generated feedback, ensuring a more reliable and equitable learning experience.

3. Content Ingestion

Content ingestion in the context of Generative AI refers to the process of collecting, processing, and storing large volumes of diverse data that will be used to train Generative AI models. These data can come from various sources, such as text documents, images, audio files, or videos, and are typically structured and formatted in a way that the Generative AI model can understand and learn from. The goal of content ingestion is to provide Generative AI with a rich, comprehensive dataset that enhances its ability to generate high-quality, relevant, and accurate outputs. Proper content ingestion ensures that the Generative AI model has access to a wide range of knowledge, which can improve its performance in answering questions and generating content.

Retrieval Augmented Generation (RAG) [13] is a technique that combines the strengths of information retrieval with Generative AI models to enhance the quality and accuracy of AI-generated content. In the case of higher education, the educational content of an online course offered by a higher education institution is ingested. In RAG, the ingestion process consists of splitting the content into chunks, generating vector embeddings, and storing the chunks and vectors in a vector database. A chunk can be, for example, a sentence, a paragraph, or a section. A vector embedding is a numerical representation of a chunk. Figure 1 shows two chunks from the content of a course, along with part of the embedding vector. At query time, the query is converted to a vector. The closest vectors to this vector are selected from the vector database. The associated chunks (the context) are combined with the query to form the user prompt for the Large Language Model (LLM).

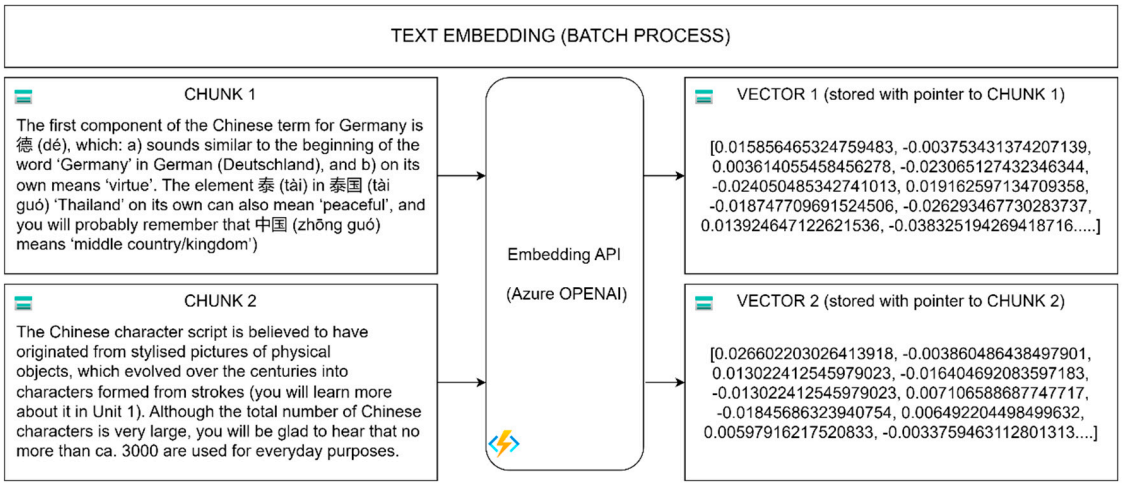

Figure 1. Chunking and text embedding.

The above approach is now referred to as Baseline RAG. However, this approach has been found to be inaccurate in some cases, as it typically only depends upon the selection of chunks using a single metric, although reranking techniques can be used to increase performance.

An improved RAG approach currently gaining popularity is GraphRAG [14]. GraphRAG allows questions to be asked of, for example, the content of several courses. There are two approaches to GraphRAG. In a Local GraphRAG approach, structured data from a Knowledge Graph is combined with unstructured data from a document. In a Global GraphRAG approach, a Large Language Model is used to extract entities and relationships from the content. These entities and relationships are stored on a graph. Clustering is then used to form communities. The text in these communities is summarised and sent along with a query to the LLM.

Typically, educational content for the learning platforms used in higher education is initially created as XML files, or using SCORM, which is based on XML [15], and then loaded into the learning platform. This has several advantages for the RAG process. The main one is that the XML tags and attributes can be stored with the ingested material, which can then be used to signpost a student to the exact location of the content. In addition, the XML can be easily converted to a Knowledge Graph for a GraphRAG approach to be adopted. A detailed overview of this architecture is shown in Figure 2.

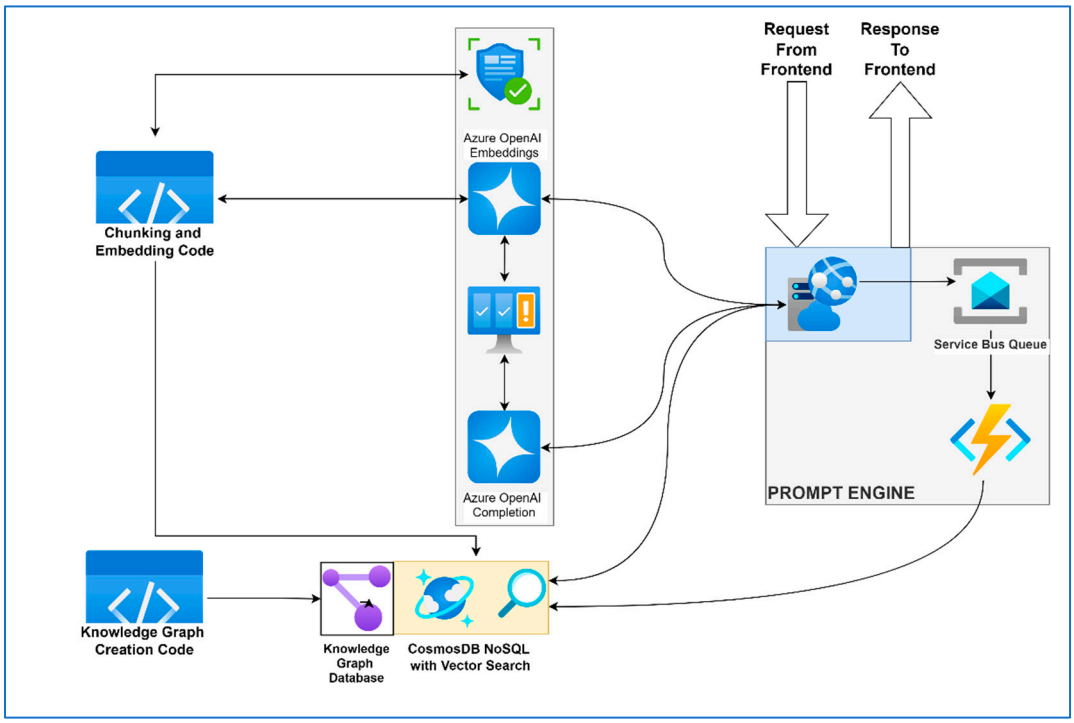

Figure 2. GraphRAG architecture for the ingestion of educational content.

In the case of the SAGE-RAI project, we have implemented four primary loaders to facilitate diverse content ingestion: *pdfLoader*, *JsonLoader*, *htmlLoader*, and *TextLoader*. This approach is designed to accommodate a broad range of file formats and data sources, thereby enhancing the system's flexibility and usability. The use of multiple content ingestion loaders addresses the growing need for systems that can efficiently process diverse data types, as noted by Stonebraker et al. [16]. The modular design of these loaders supports scalability, enabling the system to handle large-scale data processing, as discussed by Sun and Wen [17]. Additionally, this design improves the user experience by making the system more intuitive and easier to use, aligning with the principles of user-centred design [18].

4. Best Practices

The following sections outline our proposed best practices for the responsible adoption of Generative AI in higher education. These best practices collectively provide a comprehensive framework for higher education institutions to effectively and ethically integrate Generative AI into their teaching and learning practices. The aim of this framework is to enable higher education institutions to harness the potential of Generative AI while ensuring that its adoption is aligned with educational best practices and institutional goals. By focusing efforts on a responsible and human-centred approach alongside pedagogical soundness, careful planning, transparency, as well as content quality, higher education institutions can harness Generative AI's potential to enhance education while maintaining the essential human elements of teaching and learning.

4.1. Responsible Adoption

The responsible adoption of Generative AI in higher education is critical for ensuring that AI tools operate fairly, transparently, and safely. As institutions embrace Generative AI, they must address potential risks such as bias, misinformation, and violations of copyright and privacy. For example, bias is a well-documented issue in AI systems, as they are often

trained on data that reflects societal inequalities. A notable case occurred at the University of Texas at Austin, where they used AI to assist in undergraduate admissions. The AI system was initially designed to help streamline the process by identifying promising applicants based on past admissions data. However, the system began to reflect biases present in the historical data, such as favouring students from wealthier, predominantly white neighbourhoods over those from more diverse and lower-income backgrounds [19].

Generative AI systems, if not carefully managed, can perpetuate existing biases present in the data they are trained on, leading to unfair outcomes. To mitigate these risks, institutions should advocate for principles like accountability, inclusivity, and human oversight. This means establishing clear guidelines for the use of Generative AI, as well as constantly auditing Generative AI outputs for biases. Additionally, institutions must prioritise the protection of intellectual property and student privacy by implementing robust data governance policies. By adopting Generative AI responsibly, higher education institutions can harness the benefits of Generative AI technology while safeguarding against its potential harms, thereby promoting a fair and equitable learning environment for all students.

4.2. Human-Centred Approach

Adopting Generative AI in higher education should prioritise a human-centred approach, where Generative AI serves as a tool to assist educators rather than replace them. The key is to leverage Generative AI in order to enhance the capabilities of human instructors, enabling them to focus on more complex and nuanced aspects of teaching.

For instance, Generative AI can be deployed to handle routine or trivial student inquiries, such as answering frequently asked questions (FAQs) related to course logistics, assignment deadlines, or grading policies. Chatbots powered by AI can efficiently provide answers to these repetitive questions, which often consume a significant amount of time. This frees up instructors to concentrate on more critical tasks, such as providing detailed feedback on student work, facilitating deeper class discussions, or mentoring students individually on their academic growth [20]. Additionally, Generative AI tools can be used to generate draft versions of course materials or quizzes based on pre-set guidelines, which educators can then refine [10]. By automating these routine tasks, educators can redirect their energy and expertise toward fostering a richer educational experience.

By ensuring that Generative AI complements rather than competes with human educators, institutions can preserve the essential human touch in education while still benefiting from the efficiencies offered by Generative AI. This approach enhances both the administrative and pedagogical dimensions of teaching, allowing for a more personalised and effective educational experience for students.

4.3. Pedagogically Driven Solutions

Integrating Generative AI into higher education should be firmly rooted in sound pedagogical principles. The primary goal is to enhance teaching and learning experiences rather than merely adopting new technologies for their own sake. Any application of Generative AI in teaching and course development must align with established educational methodologies.

For example, in *inquiry-based learning*, where students are encouraged to ask questions, explore topics, and investigate solutions, Generative AI can provide students with access to relevant research materials, generate prompts to stimulate critical thinking, or offer feedback on their ideas. Educators can also use Generative AI to develop scenarios or case studies based on student inquiries to enable their students to explore multiple perspectives on a given issue [21]. In *constructivist learning*, where learning is seen as an active process of constructing knowledge through experience, Generative AI can assist by adapting

content to the current level of understanding of the student, or by suggesting activities or resources tailored to individual needs [22]. Additionally, in a *flipped classroom*, where students engage with learning materials independently before class, Generative AI can support the preparatory phase by answering students' questions as they work through the learning materials, by generating summaries of complex texts, or by suggesting further resources for study [23].

To ensure this alignment with educational goals, it is crucial to involve stakeholders, especially seasoned educators, in the decision-making process. Their insights and experience can help shape Generative AI initiatives that truly support these varied teaching models, making AI a genuine enhancer of learning rather than a gimmick. This approach ensures that Generative AI educational tools are not only innovative but also effective and pedagogically sound, tailored to enrich the student experience in meaningful ways.

4.4. Roadmapping

Introducing Generative AI into teaching practices within higher education institutions is a significant undertaking that requires meticulous planning. This process can be disruptive if not managed carefully, so a well-thought-out and justified approach is essential. To this end, developing a comprehensive roadmap is critical. This roadmap should outline clear objectives, define the scope of Generative AI integration, and set forth a detailed implementation plan. Key components of the roadmap include establishing milestones to track progress and performance indicators to measure the impact and effectiveness of Generative AI initiatives. By having a clear, structured plan, institutions can navigate the complexities of Generative AI adoption while minimising disruptions and maximising benefits.

4.5. Transparency

Transparency is a cornerstone of ethical Generative AI implementation in education. It is essential to clearly communicate the origins and processes involved in generating AI-driven learning materials. This includes documenting and signposting the seed data and prompts used by Generative AI to produce content. Such transparency helps build trust among students and educators, ensuring that everyone understands the source and validity of the materials they are using. By providing detailed information about the provenance of AI-generated content, institutions can foster an environment of openness and accountability, which is critical for the responsible use of Generative AI in education.

It is important that higher education institutions promote a culture of integrity and transparency, by encouraging and guiding their students to use Generative AI in a way that enhances their learning without compromising ethical standards. A good example can be seen at King's College London (https://www.kcl.ac.uk/about/strategy/learning-and-teaching/ai-guidance/student-guidance, accessed on 17 October 2024), which has issued specific guidance on how students should use Generative AI responsibly in their coursework. The guidelines emphasise that while students are encouraged to explore AI tools for learning enhancement, they must properly attribute any AI-generated assistance to their work. Similar guidance is provided by The Open University (https://about.open.ac.uk/policies-and-reports/policies-and-statements/gen-ai/generative-ai-students, accessed on 17 October 2024), which also offers training and resources to help students understand how to ethically incorporate Generative AI into their academic practices.

By applying transparency principles to both educators and students, higher education institutions can create a balanced and fair environment for the responsible adoption of Generative AI. This openness ensures that AI is used to complement, rather than undermine, the educational process.

4.6. Content Quality

The quality of learning materials remains paramount, especially when using AI-generated content in the production of learning materials. Leveraging high-quality existing learning materials in conjunction with the RAG and GraphRAG frameworks is vital for improving the accuracy and reliability of AI-generated educational content. The RAG and GraphRAG frameworks can help mitigate the risk of "hallucinations", which occur when Generative AI produces incorrect or misleading information. By ensuring that Generative AI tools are trained and augmented with robust, high-quality content, institutions can enhance the educational value of Generative AI outputs. This approach not only improves the quality of AI-generated materials but also reinforces the credibility and reliability of the educational resources provided to students.

5. Conclusions

As outlined in this paper, the responsible adoption of Generative AI in higher education offers significant potential to enhance teaching and learning practices through scalable, personalised, and high-quality educational experiences. By following the proposed best practices framework, higher education institutions can effectively integrate Generative AI in a manner that is pedagogically sound, ethically responsible, and transparent. The framework's emphasis on responsible AI ensures that the deployment of AI technologies not only enriches learning outcomes but also mitigates risks associated with misinformation, copyright infringement, and bias. Moving forward, it is crucial for higher education institutions to continuously evaluate and refine their use of Generative AI to maintain its alignment with educational goals and ethical standards. Through careful implementation and ongoing evaluation, the integration of Generative AI can potentially transform the landscape of higher education, fostering an improved learning experience for students.

Author Contributions: Conceptualization, A.M., A.E., J.K. and J.D.; writing—original draft preparation, A.M., A.E., J.K. and J.D.; writing—review and editing, A.M., A.E., J.K. and J.D.; All authors have read and agreed to the published version of the manuscript.

Funding: This work has received funding from the Engineering and Physical Sciences Research Council (EPSRC), RAI UK Award Agreement Ref: EP/Y009800/1 (SAGE-RAI).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data sharing is not applicable to this article.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Bloom, B.S. The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educ. Res.* **1984**, *13*, 4–16. [CrossRef]
2. Sabzalieva, E.; Valentini, A. *ChatGPT and Artificial Intelligence in Higher Education: Quick Start Guide*; UNESCO: Paris, France, 2023. Available online: https://unesdoc.unesco.org/ark:/48223/pf0000385146 (accessed on 17 October 2024).
3. Sharples, M. Towards social generative AI for education: Theory, practices and ethics. *Learn. Res. Pract.* **2023**, *9*, 159–167. [CrossRef]
4. Kasneci, E.; Sessler, K.; Küchemann, S.; Bannert, M.; Dementieva, D.; Fischer, F.; Gasser, U.; Groh, G.; Günnemann, S.; Hüllermeier, E.; et al. ChatGPT for good? On opportunities and challenges of large language models for education. *Learn. Individ. Differ.* **2023**, *103*, 102274. [CrossRef]
5. Chan, C.K.Y.; Hu, W. Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *Int. J. Educ. Technol. High. Educ.* **2023**, *20*, 43. [CrossRef]
6. Lee, D.; Arnold, M.; Srivastava, A.; Plastow, K.; Strelan, P.; Ploeckl, F.; Lekkas, D.; Palmer, E. The impact of generative AI on higher education learning and teaching: A study of educators' perspectives. *Comput. Educ. Artif. Intell.* **2024**, *6*, 100221. [CrossRef]

7. Bozkurt, A.; Xiao, J.; Lambert, S.; Pazurek, A.; Crompton, H.; Koseoglu, S.; Farrow, R.; Bond, M.; Nerantzi, C.; Honeychurch, S.; et al. Speculative futures on ChatGPT and generative artificial intelligence (AI): A collective reflection from the educational landscape. *Asian J. Distance Educ.* **2023**, *18*, 53–130.
8. Floridi, L. AI as agency without intelligence: On ChatGPT, large language models, and other generative models. *Philos. Technol.* **2023**, *36*, 53–130. [CrossRef]
9. Mikroyannidis, A.; Sharma, N.; Ekuban, A.; Domingue, J. Using Generative AI and ChatGPT for improving the production of distance learning materials. In Proceedings of the 24th IEEE International Conference on Advanced Learning Technologies (ICALT 2024), Nicosia, Cyprus, 1–4 July 2024.
10. Mikroyannidis, A.; Domingue, J.; Third, A.; Tarrant, D.; Pieroni, T. Smart Assessment and Guided Education with Responsible AI. In Proceedings of the 20th International Conference on Intelligent Tutoring Systems (ITS 2024), Thessaloniki, Greece, 10–13 June 2024.
11. Rienties, B.; Domingue, J.; Duttaroy, S.; Herodotou, C.; Tessarolo, F.; Whitelock, D. What distance learning students want from an AI Digital Assistant. *Distance Educ.* **2024**, 1–17. [CrossRef]
12. Ullmann, T.; Bektik, D.; Edwards, C.; Herodotou, C.; Whitelock, D. Teaching with Generative AI: Moving forward with content creation. In Proceedings of the EDEN 2024 Annual Conference, Graz, Austria, 16–18 June 2024.
13. Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.-T.; Rocktäschel, T.; et al. Retrieval-Augmented Generation for knowledge-intensive NLP tasks. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 9459–9474.
14. Edge, D.; Trinh, H.; Cheng, N.; Bradley, J.; Chao, A.; Mody, A.; Truitt, S.; Larson, J. From local to global: A Graph RAG approach to query-focused summarization. *arXiv* **2024**, arXiv:2404.16130. [CrossRef]
15. Sokol, V.; Bilova, M.; Kharin, A. An approach for creating learning content from Knowledge Management System. In *Integrated Computer Technologies in Mechanical Engineering—2020: Synergetic Engineering*; Nechyporuk, M., Pavlikov, V., Kritskiy, D., Eds.; Springer: Cham, Switzerland, 2021; pp. 692–703.
16. Stonebraker, M.; Abadi, D.; DeWitt, D.J.; Madden, S.; Paulson, E.; Pavlo, A.; Rasin, A. MapReduce and parallel DBMSs: Friends or foes? *Commun. ACM* **2010**, *53*, 64–71. [CrossRef]
17. Sun, P.; Wen, Y. Scalable architectures for big data analysis. In *Encyclopedia of Big Data Technologies*; Zomaya, A., Taheri, J., Sakr, S., Eds.; Springer: Cham, Switzerland, 2018. [CrossRef]
18. Shneiderman, B.; Plaisant, C. *Designing the User Interface: Strategies for Effective Human–Computer Interaction*, 5th ed.; Pearson: Upper Saddle River, NJ, USA, 2010.
19. Burke, L. The Death and Life of an Admissions Algorithm. 2020. Available online: https://www.insidehighered.com/admissions/article/2020/12/14/u-texas-will-stop-using-controversial-algorithm-evaluate-phd (accessed on 17 October 2024).
20. Bozkurt, A. Generative artificial intelligence (AI) powered conversational educational agents: The inevitable paradigm shift. *Asian J. Distance Educ.* **2023**, *18*, 198–204.
21. Moundridou, M.; Matzakos, N.; Doukakis, S. Generative AI tools as educators' assistants: Designing and implementing inquiry-based lesson plans. *Comput. Educ. Artif. Intell.* **2024**, *7*, 100277. [CrossRef]
22. Guo, H.; Yi, W.; Liu, K. Enhancing constructivist learning: The role of Generative AI in personalised learning experiences. In Proceedings of the 26th International Conference on Enterprise Information Systems (ICEIS 2024), Angers, France, 28–30 April 2024; Filipe, J., Śmiałek, M., Brodsky, A., Hammoudi, S., Eds.; SciTePress: Setúbal, Portugal, 2024; Volume 1, pp. 767–770. [CrossRef]
23. Wang, M.; Wang, M.; Xu, X.; Yang, L.; Cai, D.; Yin, M. Unleashing ChatGPT's power: A case study on optimizing information retrieval in flipped classrooms via prompt engineering. *IEEE Trans. Learn. Technol.* **2023**, *17*, 629–641. [CrossRef]